Collaboration Drives Individual Productivity

GORAN MURIĆ, USC Information Sciences Institute, USA

ANDRÉS ABELIUK, USC Information Sciences Institute, USA

KRISTINA LERMAN, USC Information Sciences Institute, USA

EMILIO FERRARA, USC Information Sciences Institute, USA

74

How does the number of collaborators affect individual productivity? Results of prior research have been conflicting, with some studies reporting an increase in individual productivity as the number of collaborators grows, while other studies showing that the free-rider effect skews the effort invested by individuals, making larger groups less productive. The difference between these schools of thought is substantial: if a super-scaling effect exists, as suggested by former studies, then as groups grow, their productivity will increase even faster than their size, super-linearly improving their efficiency. We address this question by studying two planetary-scale collaborative systems: GitHub and Wikipedia. By analyzing the activity of over 2 million users on these platforms, we discover that the interplay between group size and productivity exhibits complex, previously-unobserved dynamics: the productivity of smaller groups scales super-linearly with group size, but saturates at larger sizes. This effect is not an artifact of the heterogeneity of productivity: the relation between group size and productivity holds at the individual level. People tend to do more when collaborating with more people. We propose a generative model of individual productivity that captures the non-linearity in collaboration effort. The proposed model is able to explain and predict group work dynamics in GitHub and Wikipedia by capturing their maximally informative behavioral features, and it paves the way for a principled, data-driven science of collaboration.

CCS Concepts: • **Human-centered computing** → **Collaborative and social computing**; • **Software and its engineering** → *Software creation and management*.

Additional Key Words and Phrases: teamwork; collaboration; software development; productivity; Wikipedia; GitHub

ACM Reference Format:

Goran Murić, Andrés Abeliuk, Kristina Lerman, and Emilio Ferrara. 2019. Collaboration Drives Individual Productivity. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 74 (November 2019), 24 pages. <https://doi.org/10.1145/3359176>

1 INTRODUCTION

The benefits of collaboration and teamwork have been well documented. Groups can benefit from the diversity of their members [22], while the collective intelligence effects [33] enable a group to outperform even its most capable individual members [4, 15]. A group is better able to meet the multidisciplinary demands of today's increasingly more complex scientific problems, as seen in the growing shift to larger groups in science [35]. Open-source collaboration platforms exhibit a shift

Authors' addresses: Goran Murić, USC Information Sciences Institute, USA, Marina Del Rey, California, gmuric@isi.edu; Andrés Abeliuk, USC Information Sciences Institute, USA, Marina Del Rey, California, aabeliuk@isi.edu; Kristina Lerman, USC Information Sciences Institute, USA, Marina Del Rey, California, lerman@isi.edu; Emilio Ferrara, USC Information Sciences Institute, USA, Marina Del Rey, California, ferrarae@isi.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2019/11-ART74 \$15.00

<https://doi.org/10.1145/3359176>

in the distribution of work, with a few “elite” users doing most of the work [12]. There is additional evidence of the fundamental differences in how groups operate within their respective fields based by their size: small groups produce more disruptive, innovative and potentially risky work, while larger groups tend to build on the existing concepts [34].

Although groups have an advantage over individuals, it is unclear whether larger groups have an advantage over smaller groups. Even within the domain of software development, where teams of developers can vary in size up to hundreds of people, there are no clear answers to this question. Smaller teams are more agile and require less communication overhead to coordinate [7]. Smaller teams also avoid the “free-rider” effect [31], which leads individual members of larger teams to lower their effort if they perceive no impact (or proportionally lower impact) of their contributions [26]. In contrast, other research has shown that team size has a positive effect on the total amount of work invested in software developments projects [14, 23, 28]. Understanding the relationship between group size and performance can help improve productivity yielding policies to create optimal groups, e.g., by analyzing the cost-benefit of adding new collaborators versus upgrading the information technology [32].

In a quest to disentangle the relationships between group productivity and its confounding factors, some researchers have suggested that the group’s effort on a software project scales super-linearly with its size [24, 28]. This is an important finding, as it suggests that doubling a group’s size would more than double the amount of work it produces. In contrast, other studies found that productivity of software developers decreases with increasing group size [19, 25], and argue this is due to communication and coordination costs. This disagreement may be partly rooted by differences on the productivity measures, the unit of time analyzed, and the definition of a group [17]. More importantly, these previous studies only analyzed a few hundred software development groups working on highly popular software projects. In this paper, we address these issues by expanding the analysis to millions of groups across two open-source collaboration platforms.

Significance of our results

Collaboration is one of the foundations of social behavior. Understanding how the performance of groups changes as they grow is key to optimizing productivity of organizations at all scales. Our analysis of the planetary-scale GitHub and Wikipedia collaboration platforms reveals complex super-scaling dynamics for small groups, and marginal returns on productivity of increasingly larger groups. We closely examine the marginal gains of the individual efforts, and report the unexplored non-linear behaviors.

Collaboration and teamwork

Broadly speaking, a team can be defined as a group of individuals who coordinate to perform activities to achieve some common goals. The teams can differ in many ways; however, we can generally place them on a spectrum where on one side there are *structured* and on the other side there are *unstructured* teams. We can define the structured teams to have precisely defined roles, organized in advance by an external entity, with predefined communication patterns and management hierarchy. Such highly structured teams are found in traditionally organized businesses or in the military. In contrast, unstructured teams have no predefined roles and no management structure. The collaboration groups we focus on are usually somewhere in between. The emergence of the large-scale collaboration platforms, such as GitHub or Wikipedia, gave rise to a new type of *loosely structured* teams. Such teams, or collaboration groups, differ from “traditional” teams in several respects: explicit communication between collaborators may not exist; coordination is not

established in advance, but rather ad-hoc and organically; collaborators often have no predefined roles; motivations for contribution are often intrinsic. Therefore, we refer to such units as *groups of collaborators* rather than teams, even though they do exhibit some characteristics of structured teams.

2 RESULTS

Motivated by open questions and the availability of large-scale collaboration data, we address the problem of quantifying the relationship between number of collaborators and group productivity in open-source collaboration environments. We study two planetary-scale online platforms: *GitHub*, a collaboration platform, where individuals work on solo software projects or collaborate in teams, and *Wikipedia*, the World's largest encyclopedia written collaboratively by a large community of contributors. By analyzing the activity of approximately two million users and three million projects on GitHub and more than 700 thousand users collaborating on 2.6 million pages on Wikipedia, we discover that the interplay between number of collaborators and productivity exhibits complex and nonlinear dynamics, yet, both platforms display distinct commonalities. Having more collaborators is usually beneficial. We show that the marginal benefit from each additional collaborator is *super-linear*. Such benefit, however, declines as the group grows. We show that the productivity of a group as a whole is the result of increased productivity of *all individual members*, and not simply due to the presence of a few super-performing individuals. The relationship between size and productivity is not an artifact of averaging over heterogeneous group members, but holds at the individual level. When given the choice to work in multiple groups, users put more effort into projects with more collaborators. Users increase their individual activity with increasing group size, and each member is significantly more productive in a large group than when working alone or in a small group. Additionally, we analyze robustness of this finding by looking at a variety of factors which affect user's productivity.

3 RELATED WORK

Laughlin and Carey [4, 15] demonstrate that groups almost always perform better than the best individuals when solving an analytical task, both in performance and speed. Groups also tend to outperform a set of an equivalent number of independent workers [18]. Even though each member of a group puts less effort in on average, the gains of a collaboration usually dominate such losses.

For a number of tasks, besides the necessary skills, creativity plays an important role. In that case, in addition to the group size, other factors have to be considered, such as experience and time spent in a group. Creativity gets stimulated when new innovators join the group, bringing the new solutions for some old problems. However, one should keep in mind the balance of diversity in a group, as highly diversified groups could be prone to miscommunication. Besides all of the challenges, even in creative environments such as science and art, groups perform better [9]. Wuchty et al. [35] show that teams in science produce the exceptionally high-impact results. Teams also typically produce more frequently cited research than individuals do.

In general, groups have an advantage over individuals. However, it is still unclear if larger groups have an advantage over smaller groups. The answer to that question is complex and domain specific. Larger groups are more robust against external challenges [5]. Larger groups are also able to gain more success in collaborative efforts such as software development [14]. On the other hand, smaller groups are more agile which makes them more adaptable and therefore suitable for various ad-hoc projects. The benefits of small, flexible groups are particularly evident in the case of human conflicts [2]. In the area of software development, team sizes can vary from a couple of developers to hundreds of them. Having a larger team seems to have a positive impact on the project success [14], but at the same time increasing a team requires effective team coordination [7]. For example, in

Wikipedia, having more editors work on an article was associated with an increase in article quality only if the planning is done by a small subset of the contributors [13]. Both software size and team size independently have increasing returns to scale relationship with software effort [24].

Even for a single professional domain, the size of a group can affect the collective productivity in non-linear, multi-directional and conflicting ways. Increasing the group size usually has a positive effect on a total amount of work invested by the collaborators [14, 23, 28]. The underlying stimulants which drive such a trend can be explained as the motivation of an individual when collaborating with others [10]. On the other hand, group size can affect the individual and therefore total performance in the opposite way. The productivity loss occurs when individuals perceive no value to contributing or perceive no impact of their contributions to achieve the desired outcome [26]. The members of a large community tend to invest less effort in average due to the effect of *social loafing* [27]. Known also as *free-riding* [31] or *lurking* [21], it affects the communities where the effort of a fraction of individuals compensate for the rest who do not contribute. Furthermore, the late addition of new members often reflects negatively on the productivity [3] and are more likely to drop out due to social barriers [30]. A decrease in the productivity can be attributed to communication overhead [6], which becomes prominent in larger groups. If the communication is unstructured and loosely organized, the amount of information each member of a group has to process can lead to the information overload [20], which trumps the quality of communication and subsequently, the collaboration can be affected.

4 DATA AND NOTATION

4.1 GitHub

GitHub is a Web-based system for version control. It allows multiple users to work on joint projects while keeping track of the full project history and versions. GitHub allows the collaborative work of teams of various sizes on a number of project types. It is mainly oriented towards the developer community and used for facilitating the development of software, both open source and proprietary, as it supports distributed non-linear workflows. Projects on GitHub are called *repositories* and they contain the set of files related to the particular project. Besides the collaborative functions, GitHub provides a communication platform which allows developers to report issues or comment on other repositories. Users of GitHub can initiate various actions on repositories, such as *pull*, *push*, *commit*, *watch*, etc. In our analysis, we focus on users who perform *push* action. A *push* is a group of code updates uploaded to the repository. It can be performed only by users who are granted a privilege to make changes on the code. Those are the users who are trusted by the owner of the repository and make a part of the team. In our analysis, we focus only on those users considered to be the *active collaborators*. We define active collaborators as users who have an autonomy to modify the repository directly by performing a *push* action. Therefore, we did not include *pull requests* as they can be initiated by any member of a community and they are relatively rare [11]. Notice that each *push* can contain an unspecified amount of code updates. However, we use it as a proxy to identify the group of developers working together. The lines of code and the number of commits are some of the alternative measures used in literature. Previous research has shown that any of those units could be used to measure work, as the majority of pushes consist of a single commit [14] and the relation between lines of code and commits exhibits the same scaling [28].

GitHub Dataset. We used a dataset consisting of all public actions performed on GitHub in the period of 32 months starting from January 1st 2015, recorded on 1s intervals. The data used for the analysis is made of all push actions performed by users on a subset of repositories during the first three months after the repository creation. The dataset consists of 40.9 million records of *push* actions performed by 2 million users on 4.7 million repositories. Choosing the optimal time

span for analysis is a result of a trade-off between: *temporal relevance*, *contextual relevance*, and *sample size*. *Temporal relevance*: we use the most relevant data available as in average more than 90% of all events are performed during the first three months (Figure 4 in Appendix A). *Contextual relevance*: we define a group of collaborators working on a project on GitHub as a set of users actively contributing to a project over a relatively short period of time. We consider very late contributions to be less relevant to the actual group dynamic. *Sample size*: the period of three months allows the analysis of a significant number of groups of various sizes. Selecting different time windows yields similar behavior with different scales. We perform a sensitivity analysis for varying window time in Appendix B.

Additionally, we use data on user profiles and repository profiles. User profiles contain supplementary information about the user, such as a date of profile creation, the total number of created repositories, etc. Repository profiles contain information such as a date of repository creation, number of watchers, programming languages used in a repository, etc.

4.2 Wikipedia

Wikipedia is the largest and most popular online encyclopedia. The encyclopedic articles, which are referred to as *pages* are the result of collaborative effort of volunteers, who are allowed to edit entries. Any user can become an editor by visiting any page and click on “Edit”. Initially, it started as a completely open system where all authors had the same privileges and anyone could edit any page. However, certain limitations were quickly introduced to prevent undesirable edits of some particularly controversial, sensitive and/or vandalism-prone pages [1]. Every page has a corresponding *talk* page, where the editors can communicate between each other, usually by writing an additional notes about their edits. Besides the regular users, there is a number of bots which can edit pages. The bots on Wikipedia are the automated tools that carry out certain automated tasks, such as: fixing errors, updating data or reverting vandalism. In our analysis we focus on edits performed only by registered users, excluding edits performed by bots and similar scripts which have a user status, but are actually machines. Besides the number of edits, an alternative unit is the size of the edit measured by the number of bytes changed. Additional analysis described in Appendix F show that there is a similar scaling relation between number of edits and size of edits and that both of them could be used as a valid measure of contribution.

Wikipedia Dataset. The data used for this analysis is extracted from December 2018 dump of English Wikipedia edits. The dump contains all relevant revision metadata on more than 60 million pages from their creation to date. Besides the data on main articles (5.8M), the dump contains edit history of other page types such as: *talks*, *user pages*, *user talks*...

Here, we analyze a sample of main Wikipedia pages chosen uniformly at random. We analyze 2.6 million pages with 23 million edits performed by 700 thousand users. We exclude *redirect* pages, as their purpose is to automatically send visitors to another page, usually an article or section of an article. In average, Wikipedia article is 105 months old and the time difference between first and last edit is 93 months. Such a longevity allows contributions from multiple users over the period of many years. One can argue that the editors who make the updates long after each other do not collaborate. Our focus is on collaboration where people work in structured or unstructured teams, meaning that they make edits successively or alternately in relatively short period of time. In average more than 36% of all edits are performed in just three months after the page creation (Figure 4 in Appendix A) and therefore we focus our analysis only in the first three months of edits, regardless of time of creation. Similarly to GitHub, our results hold for varying time windows. More details in Appendix B.

Table 1. A short overview of terms and notations used in the paper

Term	Definition
<i>User, u</i>	An active member of GitHub or Wikipedia
<i>Project, P</i>	Either a repository on GitHub or a page in Wikipedia
<i>Number of projects, P_s</i>	Number of projects the user is contributing to
<i>Group, G</i>	A set of collaborators who work on the same project
<i>Group size, N</i>	The number of collaborators belonging to a group, $N = G $.
<i>Work, w</i>	Number of contributions made by a user. The work of a group is $W = \sum_i w_i$
<i>Effective group size, n</i>	$n = 2^H$, where $H = -\sum_{i=1}^N f_i \log_2 f_i$, and $f_i = w_i/W$
<i>Aggregate focus, F</i>	$F = \sum_{i=1}^N c_i$, and $c_i = w_i / \sum_{i,k} w_{i,k}$
<i>User age, A_u</i>	Time in days measured from the creation of user account to the end of the data range
<i>Watchers, W_c</i>	Number of watchers of a repository on GitHub
<i>Forks, F_o</i>	Number of times a repository was "forked"
<i>Repository age, A_r</i>	Time in days measured from the creation of a repository on GitHub to the last recorded push within the data range
<i>Followers, F_l</i>	Number of followers of a user on GitHub
<i>Owned repositories, O_r</i>	Number of repositories on GitHub the user created
<i>Repository description, D_{sc}</i>	Number of characters in a GitHub repository's description field
<i>Page age, A_p</i>	Time in days measured from the creation of a page on Wikipedia to the last recorded edit within the data range
<i>Largest group, N_{max}</i>	The largest group of collaborators the user is working with
<i>Smallest group, N_{min}</i>	The smallest group of collaborators the user is working with
<i>Edit size, E_s</i>	The size of an edit on Wikipedia page measured in bytes
<i>Number of created pages, P_c</i>	The number of pages on Wikipedia a user created

4.3 Notation and Definitions

Let us define some terms we use throughout this paper. A brief overview of all relevant definitions, for future reference, is summarized in Table 1.

User (u) is an active user of GitHub or Wikipedia who performed at least one *contribution* during the observed time period, and $u \in \mathbb{U}$, where $\mathbb{U}$ is a set of all users.

Project (P) is either a repository on GitHub or a page on Wikipedia.

Group (G) is a set of users $G \subset \mathbb{U}$ who have performed at least one *contribution* in the same project during the observed time period, and $G \in \mathbb{G}$, where $\mathbb{G}$ is a set of all groups. Collaborators are the users who belong to a group, where $u_i \in G$ and $i = [1 \dots N]$. Note that a single user can be a member of multiple groups at the same time. We limit our research on groups with at most 20 collaborators on GitHub, as only 0.006% of all groups have more than 20 collaborators. To maintain the statistical significance, we ignore the large groups. Small number of groups with more than 20 collaborators combined with the substantial variability can yield statistically non significant results.¹ A distribution of group sizes on Wikipedia is less skewed and we are able to analyze larger groups with up to 70 collaborators. The number of groups with more than 70 collaborators drops significantly with each additional collaborator.

Group size (N) is the number of collaborators belonging to a group as defined above, and it is simply the cardinality $N = |G|$.

Contribution is either *push* in GitHub or *edit* in Wikipedia

Work (w) quantifies the contributions of a user. The work of a group is the sum of all contributions of its members in the same project, $W = \sum_i w_i$.

¹The abundance of data allows a proper analysis of group sizes whose occurrence is relatively small, but is large enough in absolute numbers. For example, there are only 26 groups with exactly 20 members in our sample on GitHub.

Fig. 1. Average work per collaborator increases with group size. For larger groups, that trend fades out. Dotted line is a power-law fit, with an exponent $\alpha_1 = 0.28$ for GitHub and $\alpha_2 = 0.17$ for Wikipedia. It illustrates a super-linear trend suggesting that each additional collaborator increases the average group productivity. **Inset:** Distribution of group sizes has a long tail, where the groups of one (single contributors) make the overwhelming majority in both platforms.

Effective group size (n) accounts for the relative number of contributions of group members in order to correct for the possible skewed distribution of work. It is defined as $n = 2^H$, where $H = -\sum_{i=1}^N f_i \log_2 f_i$, and $f_i = w_i/W$ [14].

Aggregate focus (F) quantifies a fraction of contributions each member of a group brings to the project relative to their total contribution to all projects. It recognizes the fact that users usually work on multiple projects nearly simultaneously, allocating their time and effort among them. The contribution of a user to a project is a fraction of user's total contributions to all projects. To capture the *aggregate focus* of a group we use $F = \sum_{i=1}^N c_i$, and $c_i = w_i / \sum_{k=1}^T w_{i,k}$, where w_i is the work user i performed on a project, and $w_{i,k}$ is the work user i performed on a project k , where $k = [1, \dots, T]$. In other words, c_i is a fraction of a total user's work dedicated to a particular project.

5 GROUP PERFORMANCE AND GROUP SIZE

First, we focus on the empirical analysis of the relation between group size N (nominal number of collaborators) and group performance $\bar{W}$.

Group performance is calculated as the average amount of work per member in a group as $\bar{W} = W/N$. Figure 1 shows the group performance $\bar{W}$ as a function of group size N . Each point on the plot represents the average work per collaborator for groups aggregated by their size, with vertical bars as 95% confidence intervals. Dashed lines represent the power-law fit of the following functional form: $\bar{W} = N^\alpha \times c$, where $\alpha_1 = 0.28$, $\alpha_2 = 0.17$ and $c_1 = e^{2.45}$, $c_2 = e^{0.79}$ for GitHub and Wikipedia respectively. The positive value of the exponent α in both platforms, suggests super-linear scaling of group performance as number of collaborators grows. The total work W as a function of group size can be obtained as

$$W(N) = c \int N^\alpha dN = \frac{1}{\alpha + 1} N^{\alpha+1} + C, \quad (1)$$

which yields

$$W_1(N) = 0.781N^{1.28} + C \quad \text{for GitHub, and}$$

$$W_2(N) = 0.854N^{1.17} + C \quad \text{for Wikipedia}$$

The scaling exponent we find is consistent with previous work showing that the median exponent is $4/3$ on open source software projects [28]. To illustrate the importance of this improvement, consider a group of two collaborators on GitHub. A member of a group of size two makes on average 15 contributions, leading the group as a whole to produce 30 contributions on average. A user who belongs to a group of size ten makes on average ≈ 22 contributions, resulting in the group making about 220 contributions, on average. Although the number of collaborators grows only by a factor of five, their productivity increases eight-fold ($5^{1.28} \approx 7.85$), a super-linear increase in productivity. However, notice that such a scaling factor also indicates a decreasing marginal productivity of individual performance for increasing N .

One difficulty in interpreting these results is that the average work may not correctly reflect the distribution of work among collaborators. Instead, more work by larger groups could simply result from a few highly productive outliers skewing the average [14]. Another crucial aspect to consider is the possibility that this result is due to any confounding effect with productivity. For instance, the results may be distorted by selection bias [29]—larger, more popular projects could attract or recruit more experienced and active developers on GitHub or more specialized editors on Wikipedia. Thus, it is important to select covariates to control for selection bias.

The next section is devoted to exhaustively test the robustness of the above results. First, we identify covariates that can help reduce any biases in the results. Second, in Section 6 we exploit the fact that users contribute to multiple projects in different group sizes to perform an analysis similar to a within-subjects design. The results are consistent: individuals in average intensify their work effort when groups get larger, up to the point when such a trend starts to decline.

Correcting for confounds

We present a linear regression model of group performance to control for potential confounds among features related to the project and the group. We consider a set of intuitive features as the most obvious covariates for projects on GitHub and Wikipedia.

First, we perform a multivariate linear regression, using the Ordinary Least Square linear model. We take the average work of a group member $\bar{W}$ as a measure of group performance and a set of other available features as dependent variables. To correct for the difference in scale and estimate the relative contribution of all the variables, we scale the data so that all values fall into a range between zero and one. In order to achieve a statistically significant regression coefficient, we remove highly correlated predictors. The results of the OLS model (Table 2) suggest that the group size, even when corrected for multiple confounding variables is positively related to the group performance in both platforms. Considering the fundamental differences between GitHub and Wikipedia, the features we chose are not the same for both platforms.²

GitHub: Group size N is the most important factor explaining the group performance. Number of watchers and forks play the strongest role after group size. The number of watchers and forks are the indicators of repository popularity. More popular repositories attract more attention and subsequently more watchers and forks. With that in mind, it is reasonable to assume that more popular repositories also require more work, which is confirmed by the OLS model. Aggregate focus F quantifies the fraction of work a user dedicates to a particular project compared to other projects. Larger average aggregate focus over all group members $\bar{F}$ indicates that a particular repository is one to which users focused mostly, hence yielding high performance of a group. Some factors, however, affect the work in a negative way, such as effective group size n . It suggests the existence of highly skewed work distribution among group members. It means that usually very small fraction of users do the majority of work. If the users who belong to a group work in

²For example, a *Fork* action on GitHub, which is good proxy for popularity, has no equivalent on Wikipedia.

Table 2. Ordinary Least Squares Linear Model for group performance. Outliers (above the 95th percentile in $\bar{W}$) were filtered out. $*p < 0.005$, $R_1^2 = 0.12$ and $R_2^2 = 0.1$ for GitHub and Wikipedia respectively

GitHub				Wikipedia			
<i>variable</i>	<i>notation</i>	β_i	<i>std.err.</i>	<i>variable</i>	<i>notation</i>	β_i	<i>std.err.</i>
<i>Intercept</i>		0.0315	0.000	<i>Intercept</i>		0.1782	0.000
<i>Group size</i>	N	0.9239	0.014	<i>Group size</i>	N	0.3298	0.004
<i>Forks</i>	F_o	0.2910	0.014	<i>Page age</i>	A_p	-0.1129	0.000
<i>Watchers</i>	W_c	0.2680	0.013	<i>Avg. number of projects</i>	$\bar{G}_s$	-0.1576	0.001
<i>Repo age</i>	A_r	0.0249	0.001	<i>Avg. aggregated focus</i>	$\bar{F}$	0.3975	0.009
<i>Effective group size</i>	n	-1.0927	0.013				
<i>Avg. number of projects</i>	$\bar{G}_s$	-0.0591	0.003				
<i>Avg. aggregate focus</i>	$\bar{F}$	0.2011	0.006				

many other projects at the same time, that will negatively affect the group performance. Here, it is quantified as the average number of projects, G_s .

Wikipedia: Similarly to GitHub, group size N is an important predictor of group performance, while the average aggregate focus brings even more weight. Again, the number of projects users are working on will negatively affect the group performance.

6 INDIVIDUAL WORK AND GROUP SIZE

After the analysis of group performance, we shift our focus to individual contributions. The amount of work among collaborators is not equally distributed. Therefore, higher performance of larger groups could stem from the highly productive individuals within the groups, who may be highly motivated to work on bigger projects. Also, larger projects might attract more active contributors: more experienced developers on GitHub or more specialized editors on Wikipedia. We address these questions by exploiting the fact that a significant fraction of users work on multiple projects simultaneously. By observing how users divide their effort across multiple projects enables us to measure the effect of group size on the allocation of effort. Specifically, given user u , who works on two projects p_1 and p_2 , with different numbers of collaborators N_1 (for project p_1) and N_2 (for project p_2), where $N_2 = N_1 + 1$, and corresponding amounts of work w_1 for project p_1 and w_2 for p_2 , what is the relative work gain $\beta = w_2 - w_1$.

The first approximation for solving this problem is again the OLS model, which will show how much the mean of the dependent variable w changes given a one-unit shift in the independent variable N while holding other variables in the model constant. The results of a simple multiple linear model which quantifies the individual work are given in Table 5 in Appendix D.

One challenge in modeling how users allocate their effort across groups of different sizes is the significant individual variability among users: some are very productive regardless of group size, while others make just few contributions. Also, users differ in number of projects they are contributing to at the same time. How can we compare the work gain of two very different users, where one works on tens of projects and another one in a handful of them? The second challenge is a potential non-linearity which OLS can not capture. Simple OLS model will give a single solution for all available values of independent variable, while there can be a significant difference in linear coefficients for different ranges of N . In the following sections we address these challenges. First we implement Mixed Linear Effect model (MLE) to correct for multiple differences between users. Secondly, we increase the resolution of the MLE to address the non-linearity in the model and preserve the level of explainability and statistical significance. Notice that the set of variables used to analyze the group dynamic is extended by the additional user-specific variables.

Table 3. Results of linear mixed effects model analysis for various bins of users. For all selected bins, the coefficient which quantifies the effect of group size N is positive, suggesting that the marginal work gain for additional collaborator is positive regardless of the random variable, meaning that individual users in average put more effort in projects with more collaborators. Outliers (above the 95th percentile in $\bar{w}$) were filtered out. $*p < 0.005$

grouping by	GitHub				Wikipedia			
	group size coef: β_1	bins	max bin size	mean bin size	group size coef: β_1	bins	max bin size	mean bin size
No. of projects, G_s	2.412	164	144,464	2,387.8	0.014	1737	410,047	3,702.2
Individual work, w_j	0.980	5	221,674	78,318.8	0.005	5	1.7m	1.2m
Followers, F_l	2.425	5	235,303	78,318.8				
User ID	2.786	318,187	535	1.2	0.066	198,547	1099	3.5

Mixed Linear Effect model

Mixed Linear Effect (MLE) models are the extension of simple linear models to allow both fixed effects (to determine the impact of group size), and random effects (to allow for differences among individual users). Given that many users contribute within multiple groups of different sizes, we can perform an analysis equivalent to a within-subjects design, where one has repeated measurements of the subject at different treatment levels (group sizes) [16]. By using a Mixed Linear Effect model we can account for the productivity variability existing across different contributors and focus on the average effect of group size for contributors working in multiple groups.

In order to estimate the general trends, we perform MLE analysis by binning users by different attributes: the number of projects they work in (G_s) and the level of their individual work (w_j) for GitHub and Wikipedia. Additionally we bin for number of followers they have (F_l) for GitHub only. Our analysis shows that regardless of the eg. number of projects, users tend to put more effort into projects with more collaborators. The results in Table 3 show that the same applies to all the binning variables. The coefficients of group size obtained from the MLE models vary from 0.980 to 2.425 for GitHub and 0.005 to 0.014 for Wikipedia, indicating the positive relation between group size (N) and work (w).

The abundance of data allows for even more fine-grained binning. To correct for the largest source of variability among users, we bin users at the individual level, treating each user as an independent bin. If a user works in multiple groups of different sizes, we calculate the linear coefficient as a relation between invested work and a group size. The users who work in multiple groups of equal sizes are excluded from the analysis. The obtained coefficient $\beta_1 = 2.786$ for GitHub and $\beta_1 = 0.066$ for Wikipedia is positive, suggesting that in general users prefer to invest more work in projects with more collaborators. Binning at the individual level is the ultimate binning strategy which eliminates many problems arising when trying to put subjects in a bin based on a single characteristic. This way we are able to estimate the trend regardless of the user's baseline work level, additionally correcting for number of confounds as shown in Eq. 2 and Eq. 3. It is possible only when the number of observations is large enough to ensure the statistical significance.

The marginal gain in individual work for each additional collaborator in a group is much more significant on GitHub compared to Wikipedia. The platforms which we observe differ in many ways. Firstly, the entry "cost" is lower on Wikipedia as the amount of knowledge required to make an edit on a page is minimal. On the other hand, GitHub is mostly oriented to software engineers, and making a contribution to an existing repository requires a specific skill set. Secondly, the collaboration on GitHub requires a higher level of coordination. Editing a page on Wikipedia

Fig. 2. Marginal work gain for an additional collaborator. Each dot represents a marginal gain for additional collaborator (y -axis) depending on the range of the group size (x -axis). Vertical lines represent 95% confidence intervals. The grey area on the bottom represents the sub-linear regime, where the work gain becomes negative.

usually does not require an approval from the creator of the page, while on GitHub, the users who are allowed to make immediate change are authorized by the creator of the repository. The groups on GitHub are closer to *structured teams* compared to Wikipedia. Those and many other differences could be used to explain the discrepancy in the marginal gain of invested work for each additional group member.

6.1 Chained LME model

The previous analysis implies there is a positive association between group size and work performed in a group. However, we are interested into finding out if a similar positive relation persists over the all ranges of group sizes. Therefore we increase the resolution of the linear model and observe the change in the linear trend. To do that, we perform the MLE analysis for various ranges of group sizes. One approach is to divide the independent variable N to k ranges and then to observe the linear trend within the ranges independently. We extend this principle by selecting overlapping ranges of N . We refer to it as the *chained* linear model. Chained LME model makes up of multiple LME models built on incremental and overlapping ranges of an independent variable N . This method yields multiple coefficients for all values of N . That way, besides the prevailing linear trend we can also observe its change. How to chose the adequate width of ranges? Too narrow range would increase the resolution, but at the same time reduce the number of observations in each range. We chose the ranges to increase the resolution as much as possible and also to maintain a statistical significance at the same time. It turns out that for GitHub, a practical range width is 7 and for Wikipedia it is 10.

Again, we are primarily interested in a relationship between individual work within a group (w) and a group size (N). To correct for the most of the variability among users, each user is put in a separate bin and the data is filtered in a way to ensure that each bin consists of at least two records of different group sizes. That way we can estimate the trend for multiple group size ranges and aggregate it at the user level. Given the choice to work in groups of different sizes, the users increase the number of contributions in larger groups. We call the change in the amount of work for one unit change in group size a *marginal gain*. The marginal gain in this case is a β_1 coefficient corresponding to group size N variable in MLE model. It quantifies the relative change in the amount of work of a user for an additional collaborator.

Table 4. Results of the chained Mixed Linear Effect analysis. The obtained β_1 coefficient quantifies the relation between the group size N and individual work w . MLE model has been applied on the selected ranges of group sizes. For all selected ranges, β_1 coefficient is positive, suggesting that users, when given choice, put more effort in projects with more collaborators. Such preference becomes less prominent for larger ranges. Outliers (above the 95th percentile in $\bar{w}$) were filtered out. $*p < 0.005$

GitHub					Wikipedia				
group size range	group size coef. β_1	observ.	bins	avg. bin size	group size range	group size coef. β_1	observ.	bins	avg. bin size
1-6	4.36	173,046	59,880	2.9	1-9	0.13	176,153	16,436	10.7
2-7	4.03	106,228	36,593	2.9	4-12	0.17	172,624	16,917	10.2
3-8	3.44	41,754	15,100	2.8	7-15	0.2	172,406	18,061	9.5
4-9	2.94	20,590	7,493	2.7	10-18	0.21	172,406	20,297	8.5
5-10	2.38	11,050	4,093	2.7	13-21	0.19	172,469	22,017	7.8
6-11	1.9	6,510	2,455	2.7	16-24	0.17	118,087	17,713	6.7
7-12	1.7	4,370	1,703	2.6	19-27	0.15	80,187	13,827	5.8
8-13	1.45	2,848	1,139	2.5	22-30	0.14	58,696	11,211	5.2
9-14	1.36	1,953	797	2.5	25-33	0.12	44,558	9,221	4.8
10-15	1.43	1,287	552	2.3	28-36	0.11	35,204	7,685	4.6
11-16	1.47	851	354	2.4			...		
12-17	1.52	588	244	2.4	55-63	0.06	9,195	2,736	3.4
13-18	0.61	427	187	2.3	56-64	0.07	8,836	2,640	3.3
14-19	1.4	385	165	2.3	59-67	0.06	8,133	2,454	3.3
15-20	0.82	324	149	2.2	62-70	0.07	6,639	2,069	3.2

To ensure the validity of the results we correct for the number of possible confounding variables³. We use the Structured Sum-of-Squares Decomposition (S3D) [8] algorithm to select the most important confounding variables. The procedure for selecting the most important factors is explained in Appendix C. For GitHub we chose: Effective group size (n), Watchers (W_c), User age (A_u), Aggregated focus (F), Owned repositories (O_r), Followers (F_l) and Number of projects (G_s). The analytical form of the MLE model used for each group size range on GitHub reads as:

$$w = \beta_1 N + \beta_2 n + \beta_3 W_c + \beta_4 A_u + \beta_5 F + \beta_6 O_r + \beta_7 F_l + \beta_8 G_s + (1|userID) \quad (2)$$

For Wikipedia we chose: Effective group size (n), Number of projects (G_s), Total number of edits (W), The largest group of a user (N_{max}), Page age (A_p), Average size of user's groups ($\bar{N}$) and Number of created pages (P_c). The analytical form of the MLE model used for each group size range on Wikipedia reads as:

$$w = \beta_1 N + \beta_2 n + \beta_3 G_s + \beta_4 W + \beta_5 N_{max} + \beta_6 A_p + \beta_7 \bar{N} + \beta_8 P_c + (1|userID) \quad (3)$$

The marginal gains for GitHub and Wikipedia are shown in Figure 2, which reads as follows: on the y -axis there are various ranges of group sizes; on the x -axis there are the corresponding marginal gains for each range. For example, in case of GitHub, the user who works on a solo project is likely to contribute 4 times more if having one more collaborator. The results of LME analysis are also given in Table 4. For all selected group size ranges, the coefficient β_1 related to group size is positive, suggesting that users do invest more work in projects with more collaborators.

³Please refer to Table 1 for definitions

Fig. 3. Model prediction of work as a function of team size. The means are estimated values obtained from the chained MLE model using the iterative process. The shaded area is 95% confidence interval. Inset: The same data plotted on a log-log scale, suggesting that super-linear effect starts to fade out after $N \approx 10$

However, the values of coefficients drop with increasing group size range. It suggests that the work per group member exhibits the increasing trend followed with a saturation. Notice that for the same relative change, there can be multiple marginal gains. The gain between groups of 5 and groups of 6 collaborators can have five different gains, depending on the range we observe. We address that by building a single model by averaging the gains obtained by multiple LMEs. We explore the robustness of these results to the choice of window size in Appendix B

6.2 Unified model

Multiple linear models can give an estimation of the effect of the group size to the individual work for various ranges of group sizes. Here, we attempt to unify multiple linear models into a single non-linear model. We use the results from the chained MLE model to give a characterization of the non-linear relationship between group size and productivity of individuals.

We define the ranges of group sizes as set of tuples $R = \{(1, 1 + v), (2, 2 + v), \dots, (m, m + v)\}$, where v is a range width⁴, and m is the number of ranges used for modeling. Each range r_i has a lower bound r_i^l which is the minimal group size in a range and upper bound r_i^u , the maximal group size in a range. For each range r_i , there is a corresponding mixed effects linear model lm_i with a corresponding group size coefficient β_i . Group sizes are denoted as $k = [1 \dots N]$. Obtaining a single model from a set of linear models is an iterative process which starts from the base value wb_1 and consecutive linear models are build on top of each other. It starts with lm_1 and a base value $wb_1 = 1$, which means we start building a non-linear model for a user who makes one unit of work in a solo team. Using lm_1 we then estimate the work for all group sizes in the range r_1 . Then we move on to lm_2 , and we chose the base value as the previously estimated value for group size 2. It is followed by estimating work for all group sizes in range r_2 . In each iteration we chose the next linear model and the base value (starting point) for that model. The base value is chosen as a mean value of all previously estimated values for a starting group size. The iterative model is depicted in Figure 8. The work in a group of size k estimated using lm_i is:

$$w_k | lm_i = f(wb_i, lm_i, k), \quad r_i^l < k < r_i^u \quad (4)$$

⁴ $v = 6$ for GitHub and $v = 9$ for Wikipedia

The work in group of size k is then:

$$w_k = \overline{w_k | lm_i}, \quad \forall i \quad \text{where} \quad k \in [r_i^l \dots r_i^u] \quad (5)$$

The base work for a linear model lm_i is:

$$wb_i = \overline{w_k}, \quad k = r_i^l$$

The result after all the iterations is a set of estimated values of w (individual work) for each group size. In Figure 3, we show the means and the 95% confidence interval for the estimated work in various group sizes.

The figure illustrates the functional form describing work w invested by a member of a group as a function of the group size N . Consistent with our previous analyses we can observe that the marginal gain of having more collaborators is positive for all group sizes. It does however decrease rapidly to be negligible after certain point. Moreover, looking at the log-log plot in the inset of the figure, we have conclusive evidence that the super-linearity nature decreases beyond 10. Hence, rather than having a fix scaling exponent like in Equation 1, a better characterization of the data would be to have a decreasing scaling exponent.

7 CONCLUSION

Identifying the universal principles which apply to group productivity appears to be within the interest of various scientific disciplines. The complexity of human interactions and the heterogeneity of groups in terms of types and purposes makes this quest particularly challenging. Prior research demonstrates multiple conflicting results about the relationship between group size and productivity. In our attempt to untangle certain principles of collaborative work, we analyzed GitHub, a Web-based collaboration platform for software projects and Wikipedia, online encyclopedia. First we identify some limitations of our work and then we summarize our main findings and contributions.

We acknowledge that our results on collaborative work should not be unconditionally generalized across domains. The data used in the study is collected from two different platforms: GitHub, used mostly by people who work in software development and Wikipedia which attracts users of many various interests. Each platform imposes its norms and dictates particular behaviors which can trigger domain-specific communication and work patterns. The groups in other work setups could exhibit a different behavior, hence one should be cautious about the absolute generalization of the provided results. The definition of productivity varies among disciplines and authors. However, we equate the productivity with the amount of work, ignoring the nuances such as the time invested or the actual value of work.

We used data covering more than 40 million actions performed by 2 million users working on 4.7 million projects on GitHub and 23 million edits performed by 700 thousand users on 2.6 million pages on Wikipedia. We explore the super-linear relationship of productivity as a function of number of collaborators on a project and show that adding a new collaborator in a group is often beneficial. An additional collaborator boosts the average work in a group, suggesting a synergistic nature in groups that causes users to perform significantly better than when working individually.

We show that the productivity of a group as a whole is the result of increased productivity of all individual members, and not simply due to the presence of a few super-performing individuals. The heterogeneity and high variance of users' productivity makes it difficult to claim that the above phenomenon is more than an aggregate behavior and holds at the individual user level. However, by exploiting the fact the users work on multiple projects, we are able to demonstrate that people are more productive in larger groups. Given the choice to work on multiple projects, people put more effort into projects with more collaborators. The relation between group size and individual invested work is positive across all observed group sizes and it is particularly strong

when the groups are still relatively small. For larger groups, that relation becomes less prominent and approaches zero. Overall, our results show that there is a super-linear relationship, but it is characterized by decreasing scaling exponent as groups become larger. This shows that the high productivity in larger groups is not the artifact of a small fraction of highly-productive individuals, but rather the result of a more universal preference pattern.

Besides the group size, we analyze a set of other factors which can drive user's productivity. We use S3D, a version of regression tree algorithm, to find the most informative features. Then, we extend this analysis by quantifying the linear relationships between the invested work and corresponding confounds.

ACKNOWLEDGMENTS

The authors are grateful to the Defense Advanced Research Projects Agency (DARPA), contract W911NF-17-C-0094, for their support.

REFERENCES

- [1] [n. d.]. Wikipedia:Protection policy - Wikipedia. https://en.wikipedia.org/wiki/Wikipedia:Protection_policy
- [2] Juan Camilo Bohorquez, Sean Gourley, Alexander R. Dixon, Michael Spagat, and Neil F. Johnson. 2009. Common ecology quantifies human insurgency. *Nature* 462, 7275 (dec 2009), 911–914. <https://doi.org/10.1038/nature08631>
- [3] Frederick P. Brooks. 1995. *The mythical man-month : essays on software engineering*. Addison-Wesley Professional. 322 pages.
- [4] Harold R. Carey and Patrick R. Laughlin. 2012. Groups perform better than the best individuals on letters-to-numbers problems: Effects of induced strategies. *Group Processes & Intergroup Relations* 15, 2 (mar 2012), 231–242. <https://doi.org/10.1177/1368430211419174>
- [5] Maxime Derex, Marie-Pauline Beugin, Bernard Godelle, and Michel Raymond. 2013. Experimental evidence for the influence of group size on cultural complexity. *Nature* 503, 7476 (nov 2013), 389–391. <https://doi.org/10.1038/nature12774>
- [6] Massimiliano di Penta, Mark Harman, Giuliano Antoniol, and Fahim Qureshi. 2007. The Effect of Communication Overhead on Software Maintenance Project Staffing: a Search-Based Approach. In *2007 IEEE International Conference on Software Maintenance*. IEEE, 315–324. <https://doi.org/10.1109/ICSM.2007.4362644>
- [7] Samer Faraj and Lee Sproull. 2000. Coordinating Expertise in Software Development Teams. *Management Science* 46, 12 (dec 2000), 1554–1568. <https://doi.org/10.1287/mnsc.46.12.1554.12072>
- [8] Peter G Fennell, Zhiya Zuo, and Kristina Lerman. 2018. Predicting and Explaining Behavioral Data with Structured Feature Space Decomposition. (oct 2018). arXiv:1810.09841
- [9] Roger Guimerà, Brian Uzzi, Jarrett Spiro, and Luís A. Nunes Amaral. 2005. Team assembly mechanisms determine collaboration network structure and team performance. *Science* 308, 5722 (2005), 697–702. <https://doi.org/10.1126/science.1106340> arXiv:NIHMS150003
- [10] Guido Hertel, Norbert L. Kerr, and Lawrence A. Messé. 2000. Motivation gains in performance groups: Paradigmatic and theoretical developments on the Köhler effect. *Journal of Personality and Social Psychology* 79, 4 (2000), 580–601. <https://doi.org/10.1037/0022-3514.79.4.580>
- [11] Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M. German, and Daniela Damian. 2014. The promises and perils of mining GitHub. In *Proceedings of the 11th Working Conference on Mining Software Repositories - MSR 2014*. ACM Press, 92–101. <https://doi.org/10.1145/2597073.2597074>
- [12] Aniket Kittur, Ed Chi, Bryan A Pendleton, Bongwon Suh, and Todd Mytkowicz. [n. d.]. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. *World wide web* 1, 2 ([n. d.]), 19.
- [13] Aniket Kittur and Robert E Kraut. 2008. Harnessing the wisdom of crowds in wikipedia: quality through coordination. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*. ACM, 37–46.
- [14] Michael Klug and James P. Bagrow. 2016. Understanding the group dynamics and success of teams. *Royal Society Open Science* 3, 4 (apr 2016), 160007. <https://doi.org/10.1098/rsos.160007>
- [15] Patrick R. Laughlin, Erin C. Hatch, Jonathan S. Silver, and Lee Boh. 2006. Groups perform better than the best individuals on letters-to-numbers problems: Effects of group size. *Journal of Personality and Social Psychology* 90, 4 (2006), 644–651. <https://doi.org/10.1037/0022-3514.90.4.644>
- [16] Mary J Lindstrom and Douglas M Bates. 1988. Newton-Raphson and EM Algorithms for Linear Mixed-Effects Models for Repeated-Measures Data. *J. Amer. Statist. Assoc.* 83, 404 (1988), 1014–1022.

- [17] Thomas Maillart and Didier Sornette. 2016. Aristotle vs. Ringelmann: A response to Scholtes et al. on Superlinear Production in Open Source Software. *arXiv preprint arXiv:1608.03608* (2016).
- [18] Andrew Mao, Winter Mason, Siddharth Suri, and Duncan J. Watts. 2016. An Experimental Study of Team Size and Performance on a Complex Task. *PLOS ONE* 11, 4 (apr 2016), e0153048. <https://doi.org/10.1371/journal.pone.0153048>
- [19] Katrina D Maxwell, Luk Van Wassenhove, and Soumitra Dutta. 1996. Software development productivity of European space, military, and industrial applications. *IEEE Transactions on Software Engineering* 22, 10 (1996), 706–718.
- [20] Azadeh Nematzadeh, Giovanni Luca Ciampaglia, Yong-Yeol Ahn, and Alessandro Flammini. 2016. Information Overload in Group Communication: From Conversation to Cacophony in the Twitch Chat. (oct 2016). arXiv:1610.06497
- [21] Nigini Oliveira, Michael Muller, Nazareno Andrade, and Katharina Reinecke. 2018. The Exchange in StackExchange: Divergences between Stack Overflow and its Culturally Diverse Participants. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 130.
- [22] E Page Scott. 2007. *How the power of diversity creates better groups, firms, schools and societies*. Princeton.
- [23] Parag C. Pendharkar and James A. Rodger. 2007. An empirical study of the impact of team size on software development effort. *Information Technology and Management* 8, 4 (nov 2007), 253–262. <https://doi.org/10.1007/s10799-006-0005-3>
- [24] Parag C. Pendharkar, James A. Rodger, and Girish H. Subramanian. 2008. An empirical study of the Cobb–Douglas production function properties of software development effort. *Information and Software Technology* 50, 12 (nov 2008), 1181–1188. <https://doi.org/10.1016/J.INFSOF.2007.10.019>
- [25] Ingo Scholtes, Pavlin Mavrodiev, and Frank Schweitzer. 2016. From Aristotle to Ringelmann: a large-scale analysis of team productivity and coordination in Open Source Software projects. *Empirical Software Engineering* 21, 2 (apr 2016), 642–683. <https://doi.org/10.1007/s10664-015-9406-4>
- [26] James A. Shepperd. 1993. Productivity loss in performance groups: A motivation analysis. *Psychological Bulletin* 113, 1 (1993), 67–81. <https://doi.org/10.1037/0033-2909.113.1.67>
- [27] Yih-Chearnng Shiue, Chao-Min Chiu, and Chen-Chi Chang. 2010. Exploring and mitigating social loafing in online communities. *Computers in Human Behavior* 26, 4 (jul 2010), 768–777. <https://doi.org/10.1016/J.CHB.2010.01.014>
- [28] Didier Sornette, Thomas Maillart, and Giacomo Ghezzi. 2014. How Much Is the Whole Really More than the Sum of Its Parts? $1 + 1 = 2.5$: Superlinear Productivity in Collective Group Actions. *PLoS ONE* 9, 8 (aug 2014), e103023. <https://doi.org/10.1371/journal.pone.0103023>
- [29] Peter M Steiner, Thomas D Cook, William R Shadish, and Margaret H Clark. 2010. The importance of covariate selection in controlling for selection bias in observational studies. *Psychological methods* 15, 3 (2010), 250.
- [30] Igor Steinmacher, Marco Gerosa, Tayana U Conte, and David F Redmiles. 2019. Overcoming Social Barriers When Contributing to Open Source Software Projects. *Computer Supported Cooperative Work (CSCW)* 28, 1-2 (2019), 247–290.
- [31] James T. Strong and Rolph E. Anderson. 1990. Free-Riding in Group Projects: Control Mechanisms and Preliminary Data. *Journal of Marketing Education* 12, 2 (aug 1990), 61–67. <https://doi.org/10.1177/027347539001200208>
- [32] Hamid Tohidi and Mohammad Jafar Tarokh. 2006. Productivity outcomes of teamwork as an effect of information technology and team size. *International Journal of Production Economics* 103, 2 (2006), 610–615.
- [33] Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. 2010. Evidence for a Collective Intelligence Factor in the Performance of Human Groups. *Science* 330, 6004 (oct 2010), 686–688. <https://doi.org/10.1126/science.1193147>
- [34] Lingfei Wu, Dashun Wang, and James A. Evans. 2019. Large teams develop and small teams disrupt science and technology. *Nature* 566, 7744 (feb 2019), 378.
- [35] Stefan Wuchty, Benjamin F. Jones, and Brian Uzzi. 2007. The increasing dominance of teams in production of knowledge. *Science* 316, 5827 (2007), 1036–1039. <https://doi.org/10.1126/science.1136099> arXiv:20

APPENDIX

A ACTIVITY ON PROJECTS

Pages on Wikipedia are relatively long-lasting. In average, Wikipedia article is 105 months old and the time difference between first and last edit is 93 months. Such a longevity allows contributions from multiple users over the period of many years. On GitHub, repositories get created with increasing pace making the average age of a repository only 3.5 months and with average difference between first and last recorded action only 2.6 months. Note that our data on GitHub does not contain any events recorded before Jan 2015, which affects some of those measures. Our strategy for selecting the relevant time frame for analysis considers several factors. The time frame we chose should: encompass a significant fraction of activities; reflect the reasonable time where a collaboration is more likely to occur (the time gap of the actions should not be too large); reflect the same phase of a project development; ensure the sufficient sample size. Since more than 90% of all actions on GitHub repositories and more than 36% of all edits on Wikipedia occur in the first three months (Figure 4), we choose the time frame of three months after the project creation.

Fig. 4. Contributions over time. Plots represent the cumulative fractions of contributions on Wikipedia and GitHub over time. Plots are divided based on the total number of collaborators on a project. The vertical lines mark the end of a 3 months period after creation of repository on GitHub or page on Wikipedia.

B SENSITIVITY ANALYSIS OF TIME WINDOW

The most active period of a repository in GitHub is right after it's inception as the average activity drops over time. The same applies for pages in Wikipedia. In this paper, we focus most of our analysis on the first three months after the creation of a repository/page. However, the duration of the time window chosen after creation of repository/page can vary. In this section, we perform a sensitivity analysis on the duration t of the chosen time window T . We let t take values from 1 to 6 months for GitHub and from 1 to 44 months for Wikipedia, as the average age of Wikipedia articles allows such a wide range. For each value of $t \in \{1m \dots 6m\}$ for GitHub and $t \in \{1m \dots 44m\}$ for Wikipedia, we apply the Mixed Linear Effects Model to assess the individual work preference. First, we plot the marginal work gain for additional group member in Figure 5. For the sake of clarity, we exclude some values of t for Wikipedia. The visual inspection suggests that the analyzed phenomenon persists regardless of the time window chosen.

Fig. 5. Marginal work gain for an additional collaborator for various time windows. Each dot represents a marginal gain for additional collaborator (y -axis) depending on the range of the group size (x -axis). Vertical lines represent 95% confidence intervals. The grey area on the bottom represents the sub-linear regime, where the work gain becomes negative. Each color represents the marginal work gains calculated over different time windows.

Additionally, to better understand the effects of varying t , we build the unified model as explained in Section 6.2 for each value of t . The unified model allows to analyze the individual work invested in groups of various sizes. Multiple unified models with varying t are depicted in Figure 6. We can represent the relation between two unified models as a function of time window t as:

$$w(N, t_{i+1}) = \alpha w(N, t_i)$$

We can approximate α by analyzing the resulting unified models and get $\alpha \approx 0.95$ for Wikipedia and $\alpha \approx 1.1$ for GitHub. Note that α is not a constant, but rather a function of t which can be written as $\alpha(t_i) = \frac{t_i}{t_i+z}$ and $\lim_{t_i \rightarrow \infty} \alpha(t_i) = 1$. As the time window t_i gets larger, the difference between two consecutive models $w(N, t_i)$ and $w(N, t_{i+1})$ gets smaller.

Fig. 6. Unified models of work as a function of team size for various time windows. The means are estimated values obtained from the chained MLE model using the iterative process. The shaded area specifies the 95% confidence interval. Each color represents the unified model calculated for different time windows.

C FACTORS OF PRODUCTIVITY

In order to better understand the elements which drive user's productivity, we focus on finding the most important factors which are able to explain the majority of the variability in the target variable w . To do that, we use Structured Sum-of-Squares Decomposition (S3D) algorithm [8]. It is a variant of regression trees, that takes as input a set of features $\mathbb{X} = \{X_i\}$, $i = 1 \dots M$ and an outcome variable Y and identifies the optimal set of m important features which are able to collectively best explain the outcome. S3D performs comparable to other tree-based methods in terms of prediction. However, its biggest advantage is its transparency and the ability to visualize and explain why predictions were made.

The method partitions the values of each important feature such that Y exhibits significant variations between blocks but small variations within each block. The procedure starts with the most informative feature and it's carried out recursively by splitting the next most informative feature, and so on. It stops when either of two criteria is met: a predefined number of features is reached or the increased variability of selecting the additional feature is not significant. Recursive data decomposition allows for the creation of the parsimonious model that explains not only the relations between the features and the target variable but the relations among the features themselves.

The measure of variability used to explain the importance of features is R^2 , which takes values between zero and one. Large values of R^2 indicate a larger proportion of Y explained by the X_i . More informative variables have a larger R^2 . During the iterative process, binning each variable adds to the total R^2 and the final R^2 is obtained as the proportion of the explained sum of squares to the total sum of squares:

$$R^2 = \frac{\sum_{p \in P_{X_j}} N_p (\bar{y}_p - \bar{y})^2}{\sum_{i=1}^N (\bar{y}_i - \bar{y})^2} \quad (6)$$

Very large values of R^2 do not always imply a very good predictive power of a model, as it can be the result of overfitting. In order to make sure that our model does not overfit, we performed k-fold cross-validation and optimize in-model hyper-parameter λ . That way we can be confident that we do not overestimate the importance of certain features. Here we focus on feature importance and use the model to learn which features can explain the target variable the best. In our case, the target variable is the amount of work of a user on a project, denoted as w . We explore one set of features for GitHub: *effective team size*, *watchers*, *forks*, *group size*, *aggregate focus*, *user age*, *repository age*, *number of projects*, *largest group size*, *smallest group size*, *followers*, *owned repositories* and *repository description*; and some additional features for Wikipedia: *work of a group*, *edit size*, *page age*, *average size of groups* and *created pages*. The results obtained by the S3D model are illustrated in Figure 7. The model was able to select eight most important features in eight runs for both platforms. Adding additional features would not significantly increase R^2 .

Using S3D, we are able to identify the most important features and to build the tree-like model to estimate the work w from the set of selected features. S3D quantifies the non-linear relationships among variables, by splitting them in the optimal number of bins. This approach, at the same time, makes it challenging to optimize the system for achieving the desired value of the target variable. For instance, *effective team size* n affects the work w , but it is unclear in which direction. The influence of n to w can be positive for one range of values of n and negative for a different range. The model recognizes such relationships and applies binning accordingly. However, for the man-managed systems such as groups of people, keeping track on those subtle non-linear relations can become challenging.

Fig. 7. The amount of variation in the target variable w explained by each feature in each iterative step of the S3D algorithm, measured by R^2 . The feature that explains the most of the remaining variable at each step is marked with the black rectangle surrounding the associated bar. Notice that some features can have the same contribution at the same step of the algorithm, but are not being selected.

D MULTIPLE LINEAR MODEL FOR INDIVIDUAL WORK

To estimate how much the mean of the dependent variable w changes given a one-unit shift in the independent variable N , while holding other variables in the model constant, we use the multiple linear regression and build an OLS model. The linear coefficients which correspond to all used variables are shown in Table 5.

Table 5. Ordinary Least Squares Linear Model on user's work. Outliers (above the 95th percentile in w) were filtered out.

GitHub				Wikipedia			
variable	notation	β_i	std.err.	variable	notation	β_i	std.err.
<i>Intercept</i>		8.0244	0.036	<i>Intercept</i>		3.6464	0.004
<i>Effective group size</i>	n	-2.0002	0.029	<i>Effective group size</i>	n	-0.2526	0.001
<i>Watchers</i>	W_c	0.0123	0.000	<i>Number of projects</i>	G_s	-0.0002	4.74×10^{-7}
<i>Group size</i>	N	1.4487	0.021	<i>Number of edits</i>	W	0.0001	3.67×10^{-7}
<i>Repo age</i>	A_r	0.0056	6.5×10^{-5}	<i>Largest group</i>	N_{max}	-0.0010	3.44×10^{-7}
<i>Aggregated focus</i>	F	0.4421	0.019	<i>Created pages</i>	P_c	-0.0001	4.72×10^{-7}
<i>Number of projects</i>	G_s	0.0039	7.4×10^{-5}	<i>Page age</i>	A_p	-0.0542	0.000
<i>Followers</i>	F_l	0.0002	3.7×10^{-5}	<i>Avg. size of groups</i>	$\bar{N}$	0.0091	0.000
<i>Owned repositories</i>	O_r	-0.0005	8.1×10^{-6}	<i>Group size</i>	N	0.1164	0.000

E BUILDING UNIFIED MODEL

Here we illustrate the iterative process of building a single model from the chained MLE model explained in Section 6.2. Each subplot in Figure 8 depicts a single iteration of the process. We start with base value $wb_1 = 1$, which means that we start building a non-linear model for a user who makes one unit of work in a solo team. Using lm_1 we then estimate the work for all group sizes in the range $r1$ (upper left subplot). We move on to lm_2 , and we chose the base value as the mean of previously estimated values for group size 2. Then we estimate work for all group sizes in range $r2$ (upper row, second subplot from the left). In each iteration we chose the next linear model and the base value (starting point) for that model. The base value is always chosen as a mean value of all previously estimated values for a starting group size. The procedure is more formally defined in Section 6.2.

Fig. 8. The iterations of building the unified model from the chained MLE model

F SCALING RELATION BETWEEN ALTERNATIVE UNIT MEASURES

To measure the productivity of developers in GitHub, beside the number of *push* actions we can use alternatives such as number of commits or lines of codes. Previous work [28] has shown, on an open software development platform similar to GitHub, that the number of commits is highly correlated with lines of code. Since the vast majority of pushes consist of a single commit [14], we stay on the safe side when using *push* as the measure of productivity in GitHub. Similar applies for Wikipedia. An alternative unit measure to number of edits is the size of edit measured in bytes. Here, we test the relation between number of edits E_n and size of edits measured in bytes E_s . Figure 9 suggests slightly positive scaling relation $E_s = E_n^\beta$ with exponent $\beta \gtrsim 1.2$ and $R^2 = 0.68$. The results are similar to those reported in previous studies and show that as the number of edits grow, the size of edits does not decrease. This way we can stay confident in our choice to use the number of edits as the measure of productivity in Wikipedia.

Fig. 9. Scaling relation between edits and size of edits in Wikipedia. For Wikipedia, the scaling exponent is $\beta \gtrsim 1.2$ and $R^2 = 0.68$. The relation between number of edits E_n and size of edits E_s exhibits the similar scaling. It suggests that any of those measures could be used to measure the contribution.

Received April 2019; revised June 2019; accepted August 2019